

# Travel Report from RTAS'02

Dan Henriksson

Lund Institute of Technology

October 25, 2002

## 1. Introduction

The 8th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'02) took place September 24-27 at DoubleTree Hotel in San Jose, California. I went there together with my supervisor Karl-Erik Årzén, and the main purpose with the trip was to present our paper "Feedback Scheduling of Model Predictive Controllers", co-authored with Anton Cervin and Johan Åkesson.

## 2. The Bay Area

The conference was held in San Jose 45 miles south of San Francisco. San Jose lies in the heart of Silicon Valley and the recreation opportunities here were rather limited. The conference venue was located right between the highway and the international airport, so it was nice to have a rental car to escape with in the evenings.

On the evening of Saturday the 21st we arrived in San Francisco, where we spent the days before the conference. Sunday was spent sight-seeing in and around San Francisco. As usual it was quite misty and chilly weather in San Francisco, whereas it was considerably nicer down in San Jose.

On Monday we visited University of California at Berkeley, where we met with Professor Edward Lee and got an survey of the current status of the Ptolemy-project. We also got an introduction to Giotto, a time-triggered language for control programming.

## 3. The Conference

The focus of the RTAS conference was to bridge the gap between researchers and practitioners in the field of embedded real-time systems. A nice balance of the conference content between academia and industry was obtained with a mix of technical papers and experience papers.

The first day of the conference was devoted to a workshop on Embedded System Codesign (ESCODES), and the main conference started on Wednesday the 25th. The main conference ran in a single track with 30 minutes presentations

(27 accepted papers out of 88 submissions). This was a nice format which gave the opportunity to hear all the presentations included in the proceedings.

The conference featured a broad range of papers of which many were quite far from my area of research. Unfortunately, not many papers focused on embedded control applications. The major disappointment of the conference was the keynote speeches, but the conference also featured two interesting panel discussions, one on RT-LINUX and one on RT-JAVA. The panels featured leading experts and future directions in these two areas were discussed.

## 4. Interesting Presentations

Below I will briefly present some of the regular presentations that I found most interesting.

**”RAP: A Real-Time Communication Architecture for Large-Scale Wireless Networks”**, C. Lu, B.M. Blum, J. Stankovic, T. Abdelzaher, and T. He

*A real-time communication architecture for large-scale sensor networks was presented. These networks are typically used to perform distributed micro-sensing and control of physical environments, e.g. in surveillance systems. The paper introduced a new packet scheduling policy called velocity monotonic scheduling accounting for both time and distance constraints. This policy was shown to be very suitable for wireless sensor networks and the end-to-end deadline miss ratios in the network were reduced.*

**”Achieving End-to-end Predictability in the TAO Real-Time CORBA ORB”**, I. Pyrali, D.C. Schmidt, and R.K. Cytron

*CORBA is a middleware intended to simplify the development of distributed applications. Next-generation distributed real-time and embedded applications will have complex quality-of-service requirements, e.g. concerning latency, jitter, and dependability. The paper presented some empirical results integrating Real-Time CORBA 1.0 in the CORBA specification to provide quality-of-service guarantees. The evaluated real-time CORBA ORB is called TAO.*

**”The Aperiodic Multiprocessor Utilization Bound for Liquid Tasks”**, T. Abdelzaher, B. Andersson, J. Jonsson, V. Sharma, and M. Nguyen

*The paper presented an multiprocessor extension to a previous utilization bound for aperiodic tasks. A special task model, called the liquid task model, is assumed. This model is representative for high-performance servers (e.g. network routers and web servers) with aperiodic workloads. The schedulability bound then enables admission controllers in these systems to determine if incoming tasks can meet their deadlines. No resources will then be spent on tasks that in the end will miss their deadline.*

## **5. Conclusions**

The conference was very interesting and gave a nice in-sight into hot topics in real-time embedded systems, both from an academic and industrial point-of-view. The next RTAS will be held in Toronto, Canada in June 2003.